arXiv:1511.05860v2 [cs.IT] 21 Nov 2015

Scampi: a robust approximate message-passing framework for compressive imaging

Jean Barbier[†], Eric W. Tramel[‡], and Florent Krzakala[‡]

† Laboratoire de Théorie des Communications, Ecole Polytechnique Fédérale de Lausanne, Faculté Informatique et Communications, CH-1015, Suisse.
‡ Laboratoire de Physique Statistique CNRS UMR 8550, Université P. et M. Curie et Ecole Normale Supérieure, 24 rue Lhomond, 75005 Paris, France.

E-mail: jean.barbier@epfl.ch, eric.tramel@ens.fr, florent.krzakala@ens.fr

Abstract. Reconstruction of images from noisy linear measurements is a core problem in image processing, for which convex optimization methods based on total variation (TV) minimization have been the long-standing state-of-the-art. We present an alternative probabilistic reconstruction procedure based on approximate message-passing, Scampi, which operates in the compressive regime, where the inverse imaging problem is underdetermined. While the proposed method is related to the recently proposed GrAMPA algorithm of Borgerding, Schniter, and Rangan, we further develop the probabilistic approach to compressive imaging by introducing an expectation-maximizaiton learning of model parameters, making the Scampi robust to model uncertainties. Additionally, our numerical experiments indicate that Scampi can provide reconstruction performance superior to both GrAMPA as well as convex approaches to TV reconstruction. Finally, through exhaustive best-case experiments, we show that in many cases the maximal performance of both Scampi and convex TV can be quite close, even though the approaches are *a prori* distinct. The theoretical reasons for this correspondence remain an open question. Nevertheless, the proposed algorithm remains more practical, as it requires far less parameter tuning to perform optimally.

1. Introduction and problem setting

Over the past decade, the study of compressed sensing (CS) [1–4] has blossomed into a large field of active research in signal processing. Of particular interest in this field is the development of sparse signal reconstruction algorithms from noisy, linear measurements. While the development of generic algorithms for this problem is well advanced, there still remains much work to be done on the development of reconstruction algorithms that exploit the statistics of particular signal classes. In this work, we are interested in the class of *natural images* which has its application within compressive imaging [5]. Natural images are considered as a class of 2D signals which posses a *piecewise-continuity*, exhibiting large regions of the signal support which are constant or slowly varying interspersed with singularities of rapid transitions, or *edges*. In this work we will see how this property of piecewise-continuity can be exploited to accurately reconstruct compressively sampled images, subject to varying degrees of measurement degradation.

For the present study, we consider the reconstruction of a vectorized image signal $\mathbf{s}^* \in \mathbb{R}^N$ composed of $N = L^2$ pixels from linear measurements $\mathbf{y} \in \mathbb{R}^M$ which represent the observations of the signal projected by the measurement matrix $\mathbf{F} \in \mathbb{R}^{M \times N}$. The measurements are corrupted by an additive white Gaussian noise (AWGN) which is zero mean and of variance Δ^*,

$\boldsymbol{\xi} \sim \mathcal{N}(\boldsymbol{\xi}; \mathbf{0}_{M\times 1}, \Delta^* \mathbf{I}_{M\times M})$ where $\mathbf{I}_{a\times a}$ is the identity matrix of size $a \times a$ and $\mathbf{0}_{a\times b}$ is a matrix of size $a \times b$ whose elements are all 0. With these definitions, the undersampled observation procedure can be written as

$$\mathbf{y} = \mathbf{F}\mathbf{s}^* + \boldsymbol{\xi} \quad \Leftrightarrow \quad y_\mu = \langle \mathbf{F}_\mu, \mathbf{s}^* \rangle + \xi_\mu, \tag{1}$$

where $\langle \cdot, \cdot \rangle$ represents an inner-product and $\mathbf{F}_\mu$ is the μ^{th} row of $\mathbf{F}$. The fundamental theory of CS asserts that with a properly chosen $\mathbf{F}$, usually random, one can accurately recover $\mathbf{s}^*$ in the under-determined setting of $M < N$, by introducing *a priori* knowledge of $\mathbf{s}^*$, namely, sparsity. A sparse signal is one which, either in the cardinal basis or some other transform basis, posses only a few non-zero components. To be more precise, in an *analysis model* [6], it is assumed that a linear transform $\Psi\mathbf{s}^*$ of the signal $\mathbf{s}^*$ to be reconstructed is sparse, while $\mathbf{s}^*$ is generally not. In this case, $\mathbf{s}^*$ is sometimes refered as *cosparse*. Complementary to this definition is the *synthesis model*, in which it is directly the signal $\mathbf{s}^*$ of interest that is assumed to be sparse, i.e. Ψ is the identity. The generic reconstruction task, in the case of AWGN corrupted observations, can thus be written as a Lasso [7] problem,

$$\hat{\mathbf{s}}_{\ell_1} = \underset{\mathbf{s}\in\mathbb{R}^N}{\operatorname{argmin}} \; ||\mathbf{y} - \mathbf{F}\mathbf{s}||_2^2 + \lambda ||\Psi\mathbf{s}||_1, \tag{2}$$

where λ is a regularization parameter whose optimal selection is generally unknown *a posteriori*, and the ℓ_p norm is defined as $||\mathbf{x}||_p \triangleq (\sum_i |x_i|^p)^{1/p}$. While one might entertain the use of generic sparse reconstruction techniques for images by recovering sparse coefficients in, say, the wavelet basis, the use of gradient variation minimization has remained state-of-the-art for image reconstruction in CS. Specifically, we are interested in the application of total variation (TV) minimization [3] for image reconstruction. In this approach, one biases the optimization towards solutions which exhibit a sparse 2D discrete gradient. In the case of anisotropic TV minimization for AWGN corrupted observations, the reconstruction problem can be written as

$$\hat{\mathbf{s}}_{\text{TV}} = \underset{\mathbf{s}\in\mathbb{R}^N}{\operatorname{argmin}} \; ||\mathbf{y} - \mathbf{F}\mathbf{s}||_2^2 + \lambda \sum_{i=1}^{N} \left(|\nabla_{s_i}^{\text{H}}| + |\nabla_{s_i}^{\text{V}}| \right), \tag{3}$$

where $\nabla_{s_i}^{\text{H}}$ and $\nabla_{s_i}^{\text{V}}$ are the horizontal and vertical discrete gradients of $\mathbf{s}$, respectively, at pixel index i. This problem has been well studied from the perspective of convex optimization, with many recent developments focused on improving the rate-of-convergence and computational efficiency of algorithms to solve this problem [8]. One robust and efficient reconstruction algorithm which we utilize in this study is the TV-AL3 of [9], based on the augmented Lagrangian approach.

Rather than studying the TV minimization problem as a convex one, we will instead adopt a probabilistic framework that mimics the original TV problem, in the sense that it will try to output piecewise-continuous solutions of the system (1). The use of statistical approaches to TV minimization is not novel to this work. For instance, the 1D TV problem, and its associated phase transitions, was studied in [10] and a reconstruction method based on approximate message-passing (AMP) was first proposed. In a related vein, there also exist applications of AMP to image representation via non-local means [11] as well as AMP in conjunction with an amplitude-scale-invariant Bayes estimator [12]. Another recent work for 1D TV proposes a joint prior defined over nearest-neighbors components, characterizing the gradient with a spike-and-slab prior within AMP [13] (SS-AMP). Finally, one of the most exciting applications of the generalized AMP [14] to 2D signals has been the work of [15], which proposed the use of a cosparse analysis model. This approach, known as GrAMPA, is also general in the sense of non-AWGN i.i.d observation channels and can be used with any sparsifying basis Ψ. GrAMPA

was shown to be a state-of-the-art approach for image reconstruction in the benchmark tests of [13, 15]. Motivated by the results of both SS-AMP and GrAMPA, we present a new AMP-based algorithm for 2D CS image reconstruction particularly adapted to natural images.

Similar to GrAMPA, the proposed method will also work with the cosparse analysis model. Furthermore, we utilize the sparse non-informative parameter estimator (SNIPE) prior proposed in [15] for auxiliary gradient variables that will be introduced in the model. We term this approach "SNIPE-based cosparse analysis AMP for imaging" (Scampi). We will detail the novel modifications of our approach with respect to GrAMPA in the sequel, but they can be summarized as *i*) an expectation-maximization learning of the noise variance, which makes the algorithm more robust to channel uncertainty, and *ii*) the introduction of auxiliary variables that allow one to relax the constraints over the differences between neighboring pixels. In terms of statistical physics, our proposed approach can be seen as a "finite temperature" version of GrAMPA. This relaxation improves reconstruction accuracy for natural images, where algorithms must be more robust to violations of the piecewise-continuous assumption.

We also show that in the limit of oracle parameter selection, the performances of TV-AL3 and Scampi are comparable, and in the low-measurement regime, the differences are almost negligible. This limiting performance raises a number of interesting questions about the nature and maximal performance of TV-based CS reconstruction of natural images.

2. Proposed model and algorithm: Scampi

2.1. The cosparse TV analysis model with dual variables

Let us define the set of all edges between neighboring pixels in the image as $E \triangleq \{\nu = (i,j) : (i,j) \in \{1,\dots,N\} \times \{1,\dots,N\}, s_i \text{ neighbor of } s_j\}$. Pixel neighborhoods can be chosen arbitrarily, but to match the TV cost function, we will limit ourselves to neighborhoods corresponding to a 4-connected lattice, placing edges between horizontally and vertically co-located pixels. Next, we introduce a set of auxiliary, or *dual*, variables which exist on the edges which describe the differences between individual pixels,

$$\mathbf{d} \triangleq [(s_i - s_j) : (i,j) = \nu \in E]. \tag{4}$$

We now define the augmented signal $\mathbf{x} \triangleq [\mathbf{s}, \mathbf{d}]^\intercal$ as the concatenation of the image representation and the vector of dual variables. We can now write out an augmented system for (1) which introduces the dual variables

$$\begin{bmatrix} \mathbf{y}_{M\times 1} \\ \mathbf{0}_{|E|\times 1} \end{bmatrix} = \begin{bmatrix} \mathbf{F}_{M\times N} & \mathbf{0}_{M\times|E|} \\ \mathbf{D}_{|E|\times N} & -\mathbf{I}_{|E|\times|E|} \end{bmatrix} \times \begin{bmatrix} \mathbf{s}_{N\times 1} \\ \mathbf{d}_{|E|\times 1} \end{bmatrix} + \begin{bmatrix} \boldsymbol{\xi}_{M\times 1} \\ \boldsymbol{\zeta}_{|E|\times 1} \end{bmatrix}, \tag{5}$$

where $\mathbf{F}$ is the original measurement matrix and the dimensions of each vectors and matrices have been indicated in subscript for clarity. The entire system can be written and solved in the reduced form: $\mathring{\mathbf{y}} = \mathring{\mathbf{F}}\mathbf{x} + \mathring{\boldsymbol{\xi}}$.

$\mathbf{D}$ is a 2D finite difference operator, i.e. the concatenation of two matrices matrices which calculate both the horizontal and vertical gradients. As stated earlier, we consider for each pixel its four closest neighbors, thus $\mathbf{D} \triangleq [\mathbf{D}_\mathrm{r}, \mathbf{D}_\mathrm{d}]^\intercal$ is the concatenation of $\mathbf{D}_\mathrm{r}$ which is made of zeros everywhere except for a main diagonal of 1's, and a single -1 on each row, in a column corresponding to the *right* neighbor of the pixel at the index represented by the 1 on the main diagonal. The matrix $\mathbf{D}_\mathrm{d}$ is constructed in similar fashion, but for the *down* nieghbor. This choice of right/down differencing, versus left/up, is arbitrary. As long as $\mathbf{D}$ is constructed consistently, all pairwise interactions are encoded consistently. Finally, the vector $\boldsymbol{\zeta}$ is an error that can be thought of as the degree of correspondence between $\mathbf{d}$ and the true discrete 2D gradient of $\mathbf{s}$. We will discuss this point more throughly in the next section. As a whole, we refer to this augmented system as a *cosparse TV analysis model with dual variables*.

2.2. Probabilistic model for TV reconstruction

We now present a Bayesian framework for the TV problem. The above construction allows one to easily associate a factorized prior P_0 to $\mathbf{x}$, a key assumption for the AMP algorithm [16–21]. While some arbitrary *i.i.d.* distribution can be assigned to the elements of $\mathbf{s}$, we will assume a SNIPE prior P_{S} [15] for $\mathbf{d}$. This prior is defined as the limiting distribution of a spike-and-slab distribution with infinite variance,

$$P_{\mathrm{S}}(x;\omega) \triangleq \lim_{\sigma\to\infty} \frac{\rho(\sigma,\omega)}{\sigma}\mathcal{N}(x/\sigma;0,\sigma^2) + (1-\rho(\sigma,\omega))\delta(x), \tag{6}$$

for which, under a proper scaling of $\rho(\sigma,\omega)$, has ω as the only free parameter of the prior. For further details on SNIPE, we direct the reader to [15]. By using P_{S} to bias the elements of $\mathbf{d}$ during inference, we place more probabilistic weight on images whose finite differences are sparse, according to ω.

If we consider that the elements of $\mathbf{s}$ obey a uniform distribution, and if we enforce that the elements of the concatenated signal $\mathbf{x}$ are ordered such that the first N elements correspond to pixel values and the last $|E|$ elements correspond to $\mathbf{d}$, then the factorized prior for $\mathbf{x}$ can be writen as

$$P_0(\mathbf{x}) \propto \prod_{i=1}^{N+|E|} P_0^i(x_i) = \prod_{i=1}^{N+|E|} \left[\mathbb{I}(i\leq N)\mathcal{U}(x_i) + \mathbb{I}(i>N)P_{\mathrm{S}}(x_i;\omega)\right], \tag{7}$$

where $\mathcal{U}$ is a uniform distribution over some domain of $\mathbb{R}$ and $\mathbb{I}(\cdot)$ is the indicator function.

With the prior defined for $\mathbf{x}$, we would like to turn our attention to the posterior distribution of $\mathbf{x}$ given $\mathbf{y}$, however, first, we must take a moment to consider the likelihood induced by the distributions of the error terms $\boldsymbol{\xi}$ and $\boldsymbol{\zeta}$. We diverge from the GrAMPA approach, by the introduction of the "finite temperature" noted in the previous section. In order to mimic the GrAMPA approach thanks to the cosparse TV analysis model with dual variables (5), one would set $\boldsymbol{\zeta} = \mathbf{0}_{|E|\times 1}$, creating a model for which the elements of $\mathbf{d}$ are exactly the discrete gradient of $\mathbf{s}$. Indeed, this is an intuitive construction. However, imposing such a strict condition might impede the reconstruction. The fist key feature of our approach is that instead, we propose that the hard linear constraints $\{d_\nu - (s_i - s_j) = 0 : (i,j) = \nu \in E\}$ be "relaxed" by taking $\boldsymbol{\zeta} \sim \mathcal{N}(\boldsymbol{\zeta}; \mathbf{0}_{|E|\times 1}, T\ \mathbf{I}_{|E|\times|E|})$. Here, the amount of relaxation is determined by the temperature-like variance parameter, T. In the frozen zero-temperature limit, $T \to 0$, one recovers GrAMPA, up to the nuances of implementation. However, for finite temperature, some slack is introduced into the problem, allowing the dual variables to fluctuate some distance away from the true gradient, a helpful feature to avoid poor local solutions.

Finally, we can now write the posterior distribution for our TV model (5) as

$$P_{\mathrm{Scampi}}(\mathbf{x} = [\mathbf{s},\mathbf{d}] | \mathring{\mathbf{y}} = [\mathbf{y},\mathbf{0}]) = \frac{1}{\mathcal{Z}} \prod_{\mu=1}^{M} \mathcal{N}\left(y_\mu; \langle \mathbf{F}_\mu, \mathbf{s}\rangle, \Delta\right) \prod_{\nu=1}^{|E|} \left[\mathcal{N}\left(d_\nu; \langle \mathbf{D}_\nu, \mathbf{s}\rangle, T\right) P_{\mathrm{S}}(d_\nu;\omega)\right], \tag{8}$$

where the normalization $\mathcal{Z} \triangleq Z(\mathbf{y},\mathbf{F},\mathbf{D},\Delta,T,\omega)$, also known as the partition function in physics, is a function of the problem instance, and is intractable to calculate, in practice. If known, the noise variance associated with the measurements in the posterior should match the true noise variance, $\Delta = \Delta^*$. If Δ^* is unknown, it may be estimated, as we will see in Sec. 2.4. Rather than attempting to calculate this posterior exactly, the message-passing approach we discuss, here, attempt to *approximate* this posterior via information transfer between nodes on the factor graph representation of the posterior. Fig. 1 gives the factor graph representation of the proposed Scampi posterior. The factor graph associated with the posterior sampled by GrAMPA is the same, but lacks the red variables and factors. Also, the

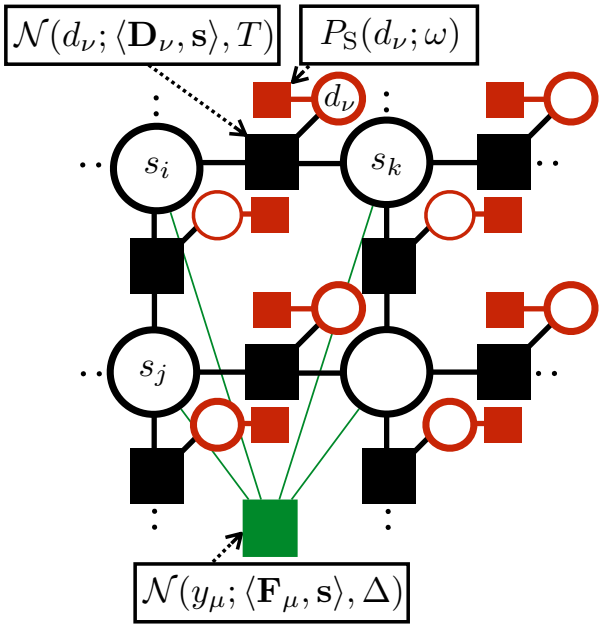

Figure 1. Factor graph associated to the probability measure (8) sampled by the Scampi algorithm for the image reconstruction. The large black circles are the pixels of $\mathbf{s}$, while the smaller red ones are the dual variables, $\mathbf{d}$. The squares represent the factors: the green ones connected to all the pixel variables are the "measurement" factors enforcing $\mathbf{s}$ to verify the linear measurements $\mathbf{y}$, up to Gaussian fluctuations of variance Δ. The red factors represent the factorizable prior we assume about the dual variables, which is the SNIPE prior P_{S}. Finally, the black factors are enforcing consistency between $\mathbf{s}$ and $\mathbf{d}$, as weighted by T. To avoid double counting, for each pixel, only two out of the four possible neighbor interactions are taken into account.

black factors would be of the form $P_{\mathrm{S}}(\langle \mathbf{D}_\nu, \mathbf{s} \rangle; \omega)$ instead of the proposed Gaussian densities, i.e. $P_{\mathrm{GrAMPA}}(\mathbf{s}) \propto \prod_{\mu=1}^{M} \mathcal{N}\left(y_\mu; \langle \mathbf{F}_\mu, \mathbf{s} \rangle, \Delta\right) \prod_{\nu=1}^{|E|} P_{\mathrm{S}}(\langle \mathbf{D}_\nu, \mathbf{s} \rangle; \omega)$.

Lastly, we would like to mention that one can also consider a 2D extension of SS-AMP without appealing to the augmented analysis model (5). In this case, the posterior becomes $P_{\mathrm{SS-AMP}}(\mathbf{s}) \propto \prod_{\mu=1}^{M} \mathcal{N}\left(y_\mu; \langle \mathbf{F}_\mu, \mathbf{s} \rangle, \Delta\right) \prod_{\nu=1}^{|E|} [\rho \mathcal{N}(\langle \mathbf{D}_\nu, \mathbf{s} \rangle; 0, T) + (1-\rho)\delta(\langle \mathbf{D}_\nu, \mathbf{s} \rangle)]$. However, empirical results indicate that this algorithm suffers from strong convergence issues. Additionally, the exploding complexity, and quantity, of the resulting message-passing equations create too many numerical issues for a practical implementation at scale.

2.3. AMP for the cosparse analysis model with dual variables: Scampi

The Scampi algorithm, detailed in Fig. 2, is canonical AMP applied to the cosparse analysis model with dual variables (5), with a few added features we detail in the sequel. Scampi performs a minimum mean square error estimation of $\mathbf{x}$ by sampling from an iteratively computed approximation of the distribution (8). The full derivation of AMP in similar notations can be found in [20, 22]. Additionally, we define the augmented vector of per-factor noise variances as $\mathring{\Delta}_\mu \triangleq \Delta \mathbb{I}(\mu \leq M) + T \mathbb{I}(\mu > M)$.

Despite the fact that the augmented system matrix $\mathring{\mathbf{F}}$ is sparse, and that the derivation of AMP is based on dense matrices, nothing prevents us from applying it effectively to such a problem, and as we will see, it achieves state-of-the-art performance. The prior-dependent

$t \leftarrow 0$
$\tau \leftarrow 1 + \epsilon$
while $t < t_{max}$ **and** $\tau > \epsilon$ **do**
  $\tilde{\Theta}_\mu^{(t+1)} \leftarrow \sum_{i=1}^{N+|E|} \mathring{F}_{\mu i}^2 v_i^{(t)}$
  $w_\mu^{(t+1)} \leftarrow \beta w_\mu^{(t)} + (1-\beta)\left(\sum_{i=1}^{N+|E|} \mathring{F}_{\mu i} a_i^{(t)} - \tilde{\Theta}_\mu^{(t+1)} \frac{\mathring{y}_\mu - w_\mu^{(t)}}{\mathring{\Delta}_\mu^{(t)} + \Theta_\mu^{(t)}}\right)$
  $\Theta_\mu^{(t+1)} \leftarrow \beta \Theta_\mu^{(t)} + (1-\beta)\tilde{\Theta}_\mu^{(t+1)}$
  $\Sigma_i^{(t+1)} \leftarrow \left(\sum_{\mu=1}^{M+|E|} \frac{\mathring{F}_{\mu i}^2}{\mathring{\Delta}_\mu^{(t)} + \Theta_\mu^{(t+1)}}\right)^{-\frac{1}{2}}$
  $R_i^{(t+1)} \leftarrow a_i^{(t)} + \left(\Sigma_i^{(t+1)}\right)^2 \sum_{\mu=1}^{M+|E|} \mathring{F}_{\mu i} \frac{\mathring{y}_\mu - w_\mu^{(t+1)}}{\mathring{\Delta}_\mu^{(t)} + \Theta_\mu^{(t+1)}}$
  $v_i^{(t+1)} \leftarrow f_{\mathrm{c}}\left(\left(\Sigma_i^{(t+1)}\right)^2, R_i^{(t+1)}; \omega\right)$
  $a_i^{(t+1)} \leftarrow f_{\mathrm{a}}\left(\left(\Sigma_i^{(t+1)}\right)^2, R_i^{(t+1)}; \omega\right)$
  Update all $\mathring{\Delta}_\mu^{(t+1)}$ according to (15)
  $\tau \leftarrow \frac{1}{N}||\mathbf{a}^{(t+1)} - \mathbf{a}^{(t)}||_2^2$
  $t \leftarrow t + 1$
end while
return $\mathbf{a}^{(t)}$

Figure 2. The Scampi algorithm with damping (controlled by $0 \leq \beta < 1$) for the cosparse model with dual variables. The convergence tolerance is given by ϵ and t_{max} is the maximum number of iterations. A suitable initialization for the posterior mean and variance are, respectively, $a_i^{(0)} = \mathbb{E}_{P_0^i}(x_i)$ and $v_i^{(0)} = \mathrm{Var}_{P_0^i}(x_i)$, where P_0^i is the prior associated with component x_i. Additionally, one initializes $w_\mu^{(0)} = \mathring{y}_\mu$. Once the algorithm has converged, the minimum mean square error estimate of component x_i is given by $a_i^{(t)}$.

denoisers which iteratively produce the posterior estimates are given by

$$\begin{aligned} f_{\mathrm{a}}(\Sigma_i^2, R_i; \omega) &\triangleq \int \mathrm{d}x\, x\, P_0^i(x)\, \mathcal{N}\left(x; R_i, \Sigma_i^2\right) \\ &= R_i\, \mathbb{I}(i \leq N) + f_{\mathrm{a}}^{\mathrm{S}}(\Sigma_i^2, R_i; \omega)\, \mathbb{I}(i > N), \qquad (9) \\ f_{\mathrm{c}}(\Sigma_i^2, R_i; \omega) &\triangleq \int \mathrm{d}x\, x^2\, P_0^i(x)\, \mathcal{N}\left(x; R_i, \Sigma_i^2\right) - f_{\mathrm{a}}(\Sigma_i^2, R_i; \omega)^2 \\ &= \Sigma_i^2\, \mathbb{I}(i \leq N) + f_{\mathrm{c}}^{\mathrm{S}}(\Sigma_i^2, R_i; \omega)\, \mathbb{I}(i > N) \qquad (10) \end{aligned}$$

where the prior terms are given by (7) and the SNIPE-specific denoisers [15] for the dual variables are given by

$$f_{\mathrm{a}}^{\mathrm{S}}(\Sigma^2, R; \omega) \triangleq \frac{R}{1 + e^{-\frac{R^2}{2\Sigma^2} + \omega}}, \quad f_{\mathrm{c}}^{\mathrm{S}}(\Sigma^2, R; \omega) \triangleq \frac{1}{1 + e^{-\frac{R^2}{2\Sigma^2} + \omega}} \left(\frac{R^2}{1 + e^{\frac{R^2}{2\Sigma^2} - \omega}} + \Sigma^2\right). \qquad (11)$$

2.4. Learning the noise variance via the Bethe free energy

We now present the second key feature of Scampi, an iterative estimation of both the noise variance and dual parameter relaxation, which improves reconstruction accuracy for natural images. However, we note that this approach in fact *degrades* reconstruction accuracy for

strictly piecewise constant signals. Nevertheless, if the image of interest is *known* to be strictly piecewise constant, the temperature T can be fixed to 0 to improve performance. If not, the hyperparameters $\{\Delta, T\}$ can be learned via the expectation-maximization we now detail.

Given the definition of the Bethe free energy G derived in [23] for the posterior measure (8), we update $\{\Delta, T\}$ at each iteration according to their fixed-point values in G. For our posterior, the Bethe free energy is written as

$$G = \frac{1}{2}\sum_{\mu}\left[\frac{(\mathring{y}_\mu - \sum_i \mathring{F}_{\mu i}a_i)^2}{\mathring{\Delta}_\mu} + \log\left(1 + \frac{\sum_i \mathring{F}^2_{\mu i}v_i}{\mathring{\Delta}_\mu}\right) + \log\left(2\pi\mathring{\Delta}_\mu\right)\right] + \sum_i D_{\mathrm{KL}}(P_i||P_0^i), \quad (12)$$

where $D_{\mathrm{KL}}(P_i||P_0)$ is the Kullback-Leibler divergence between the prior P_0^i of the i^{th} component and its estimated posterior $P_i(x_i|\Sigma_i^2, R_i) \propto P_0^i(x_i)\mathcal{N}(x_i; R_i, \Sigma_i^2)$, and the Gaussian fields $\left\{R_i, \Sigma_i^2\right\}$ and posterior mean and variance estimates $\{a_i, v_i\}$ of x_i are the fixed-point values of the Scampi algorithm in Fig. 2.

For the moment, let us consider that there is a unique noise parameter $\mathring{\Delta}$ for all of our factors. To minimize G, we start by isolating the terms which depend on $\mathring{\Delta}$,

$$G_{\mathring{\Delta}} = \frac{1}{2}\sum_{\mu=1}^{M+|E|}\left[\frac{(\mathring{y}_\mu - \sum_{i=1}^{N+|E|}\mathring{F}_{\mu i}a_i)^2}{\mathring{\Delta}} + \log\left(\mathring{\Delta} + \sum_{i=1}^{N+|E|}\mathring{F}^2_{\mu i}v_i\right)\right]. \quad (13)$$

We then write a fixed-point condition which is dependent on $\mathring{\Delta}$, for all other variables fixed,

$$\frac{\partial G_{\mathring{\Delta}}}{\partial\mathring{\Delta}} = -\frac{1}{2}\sum_{\mu=1}^{M+|E|}\left[\frac{1}{\mathring{\Delta}^2}\left(\mathring{y}_\mu - \sum_{i=1}^{N+|E|}\mathring{F}_{\mu i}a_i\right)^2 - \left(\mathring{\Delta} + \sum_{i=1}^{N+|E|}\mathring{F}^2_{\mu i}v_i\right)^{-1}\right] = 0. \quad (14)$$

It is possible to extract the optimal $\mathring{\Delta}$ by equating the two terms inside the sum, providing the set of solutions $\mathring{\Delta}_\mu \ \forall\ \mu \in \{1, \ldots, M + |E|\}$. Defining the auxiliary term $\chi_\mu \triangleq \mathring{y}_\mu - \sum_i \mathring{F}_{\mu i}a_i$, each $\mathring{\Delta}_\mu$ is obtained by solving the quadratic equation $\mathring{\Delta}_\mu^2 - \chi_\mu^2\mathring{\Delta}_\mu - \chi_\mu^2\sum_i \mathring{F}^2_{\mu i}v_i = 0$. Since all variance terms are strictly positive, the feasible solution for this equation is simply $\mathring{\Delta}_\mu = g_\mu$ where $g_\mu \triangleq \frac{1}{2}\big(\chi_\mu^2 + \chi_\mu\sqrt{\chi_\mu^2 + 4\sum_i \mathring{F}^2_{\mu i}v_i}\big)$. Averaging over $\{\mathring{\Delta}_\mu : \mu \in \{1, \ldots, M+|E|\}\}$, we obtain a single parameter $\mathring{\Delta}$. Recall, however, that we wish to consider two *different* noise variances $\{\Delta, T\}$. Since both $\mathbf{s}$ and $\mathbf{d}$ exist in different domains, it is not reasonable to expect that they both should share the same noise variance estimate. Instead, two averages are performed over the proper sets of variables in $\mathbf{x}$. Introducing the iterative time index, we obtain the fixed-point estimation

$$\Delta^{(t+1)} = \frac{1}{M}\sum_{\mu=1}^{M} g_\mu^{(t)}, \quad T^{(t+1)} = \frac{1}{|E|}\sum_{\mu=M+1}^{M+|E|} g_\mu^{(t)}, \quad (15)$$

where $g_\mu^{(t)}$ is a function of $\left\{\Delta^{(t)}, T^{(t)}, \mathbf{a}^{(t+1)}, \mathbf{v}^{(t+1)}, \mathring{\mathbf{y}}, \mathring{\mathbf{F}}\right\}$. In the final implementation of Scampi, we also introduce a damping on the update of $\{\Delta, T\}$.

Intuitively, this iterative estimation of the noise variances operates as a kind of annealing which plays a pivotal role in providing Scampi's superior reconstruction performance for natural images. Additionally, the estimation of $\{\Delta, T\}$ means that Scampi is more robust to uncertainty on these hyperparameters than GrAMPA, an important feature for practical use cases. We present results demonstrating this effect in Sec. 3. Our Scampi implementation can be downloaded on `https://github.com/jeanbarbier/scampi`.

2.5. Hadamard operators for large-scale compressive imaging

We now turn our attention to one important practical consideration for high-dimensional signals, such as natural images. As N increases, the construction, storage, and numerical use of random, dense projection matrices containing $O(N^2)$ quickly becomes computationally infeasible. To confront this issue, we use sub-sampled Hadamard operators, which have been empirically shown to provide performance very close to purely random matrices [24], but at a significant reduction in computational complexity, requiring $O(N \log N)$ operations for matrix multiplication.

Since each Hadamard mode, except the average mode, have exactly zero mean, the system (5) would be invariant by a constant shift in the signal components without the presence of the average mode in $\mathbf{F}$. To break this symmetry, we force the retention of the average mode, which fixes the mean of the signal, when randomly selecting Hadamard modes for the construction of $\mathbf{F}$. Once a realization of $\mathbf{F}$ is chosen (as well as $\mathbf{F}^\intercal$ and $\mathbf{F}^2, (\mathbf{F}^2)^\intercal$ where the squared operators are simply sums in the case of Hadamard matrices), it is trivial to construct $\mathring{\mathbf{F}}$ (and $\mathring{\mathbf{F}}^\intercal, \mathring{\mathbf{F}}^2, (\mathring{\mathbf{F}}^2)^\intercal$, as well). Furthermore, $\mathring{\mathbf{F}}$ does not generate memory issues since all sub-matrices, other than $\mathbf{F}$, are extremely sparse with complexity of $O(N)$ per matrix multiplication. Thus, the overall complexity is $O(N \log N)$ per Scampi iteration.

3. Experimental results

We now present the results of two experimental protocols for image reconstruction from AWGN corrupted CS measurements. In all cases, the tested images are rescaled to the range $s_i \in [0, 1]$, though this property isn't leveraged by any of the tested algorithms. Reconstruction performance is quantified using the normalized signal to noise ratio (NSNR), measured in decibels, between the true image $\mathbf{s}^*$ and the reconstruction $\hat{\mathbf{s}}$ defined as $\mathrm{NSNR}(\hat{\mathbf{s}}, \mathbf{s}^*) \triangleq 10 \log_{10}(||\mathbf{s}^* - \hat{\mathbf{s}}||_2^2 / ||\mathbf{s}^*||_2^2)$ dB. For all experiments, the level of the true experimental noise variance, Δ^*, is reported as the *input* SNR (ISNR) and calcualted as $\mathrm{ISNR} = 10 \log_{10}(||\mathbf{y}||_2^2 / (M\Delta^*))$ dB.

We conduct our experiments for the common test images *Lena*, *Barbara*, *Peppers*, as well as the Shepp-Logan phantom (*Phantom*) at a resolution of 512×512 pixels. *Barbara* and *Lena* represent common examples of complex natural images, with *Lena* possessing a mixture of both smooth and high-variation regions, while *Barbara* contains very many edges due to the large amount of textured content in the image. *Peppers* represents a simpler case of a more piecewise-continuous natural image, and *Phantom* is, simply, a cartoon image of a few regions of *constant* pixel value.

The first experiment, which results are presented in Fig. 3, compares the robustness of GrAMPA and Scampi, which includes the noise learning (15), to uncertainty of the measurement AWGN noise variance Δ^*. To simulate this uncertainty, both GrAMPA and Scampi are provided with an initial estimate of the noise variance, $\Delta = \gamma \Delta^*$, where $\gamma = 1$ represents perfect knowledge about the noisy channel. Each plot corresponds to the achieved reconstruction performance for a different image and two sub-sampling rates $\alpha = M/N \in \{0.1, 0.5\}$, as a function of the mismatch factor $\gamma \in [10^{-4}, 10^4]$. The dashed and solid curves correspond to the best GrAMPA and Scampi performances, respectively, over the range of the parameter $\omega \in \{0, 1, \ldots, 5\}$ for the SNIPE prior $P_{\mathrm{S}}(\cdot; \omega)$. Generally, the best ω value for GrAMPA is the same (or close) to the one for Scampi. As seen in these charts, Scampi is able to provide better NSNR performance than GrAMPA, and without requiring any specific knowledge about Δ^*. GrAMPA, however, requires that Δ^* be known, in most cases, within an order of magnitude. Additionally, in the case of images, such as *Phantom*, for high ISNR, exact knowledge of Δ^* is in fact *detrimental* to the performance of the algorithm.

In the second experiment, see Fig. 4, we show the best-case performance of both GrAMPA and Scampi as a function of the ISNR for the specified test images at $\alpha = 0.1$ and $\alpha = 0.5$. We compare both approaches to TV-AL3 [9], as well, to give a reference point between these two probabilistic techniques and a state-of-the-art convex technique. For these experiments, best-

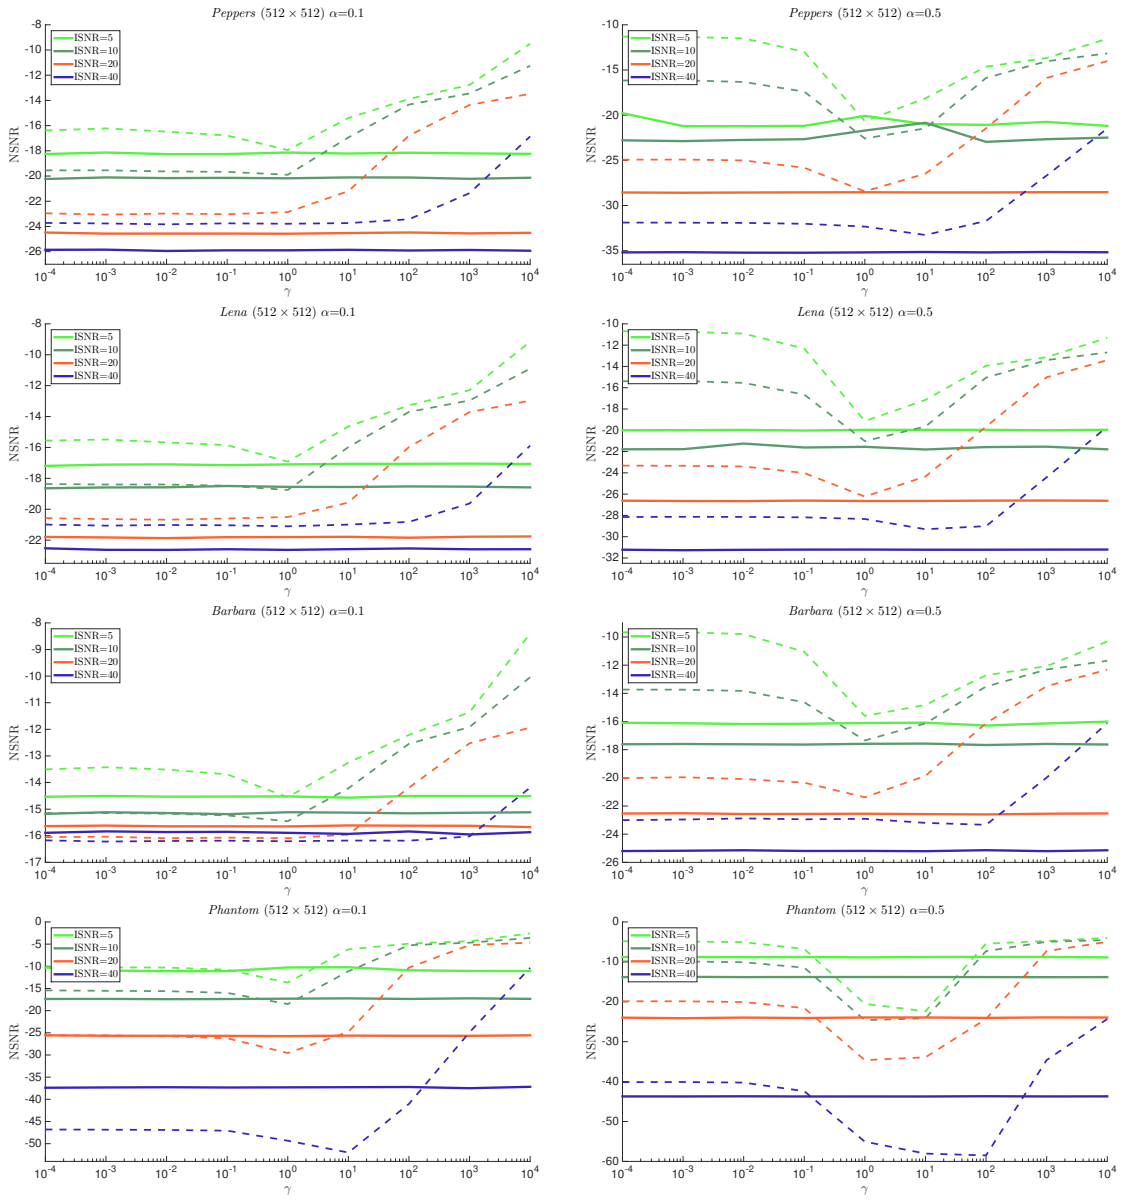

Figure 3. Reconstruction performance in NSNR (the lower the better) of Scampi with noise learning (solid curves) and GrAMPA (dashed curves), over the noise variance mismatch factor γ, for varying 512×512 test images at $\alpha = 0.1$ (left) and $\alpha = 0.5$ (right). Colors correspond to various levels of ISNR (solid and dashed curves of same color correspond to same dB). Each point is obtained through a single run for a random $(\mathbf{F}, \boldsymbol{\xi})$ instance common to both algorithms. At this problem dimensionality, finite size performance fluctuations are negligible.

case means that we intensively test over the free parameters of the algorithms, reporting only the best results. In the case of GrAMPA and Scampi, we test over ω, and for TV-AL3 we test over the weight of the TV regularization cost. In all cases, each point on the curves represents the best outcome over 30 reconstructions sweeping over each tuning parameter's domain. Additionally, we provide all algorithms with the true Δ^*. Our goal in reporting these best-case results is to compare the *maximal* performance of these techniques.

The findings we report are quite interesting and varied. The first topic we address is that for

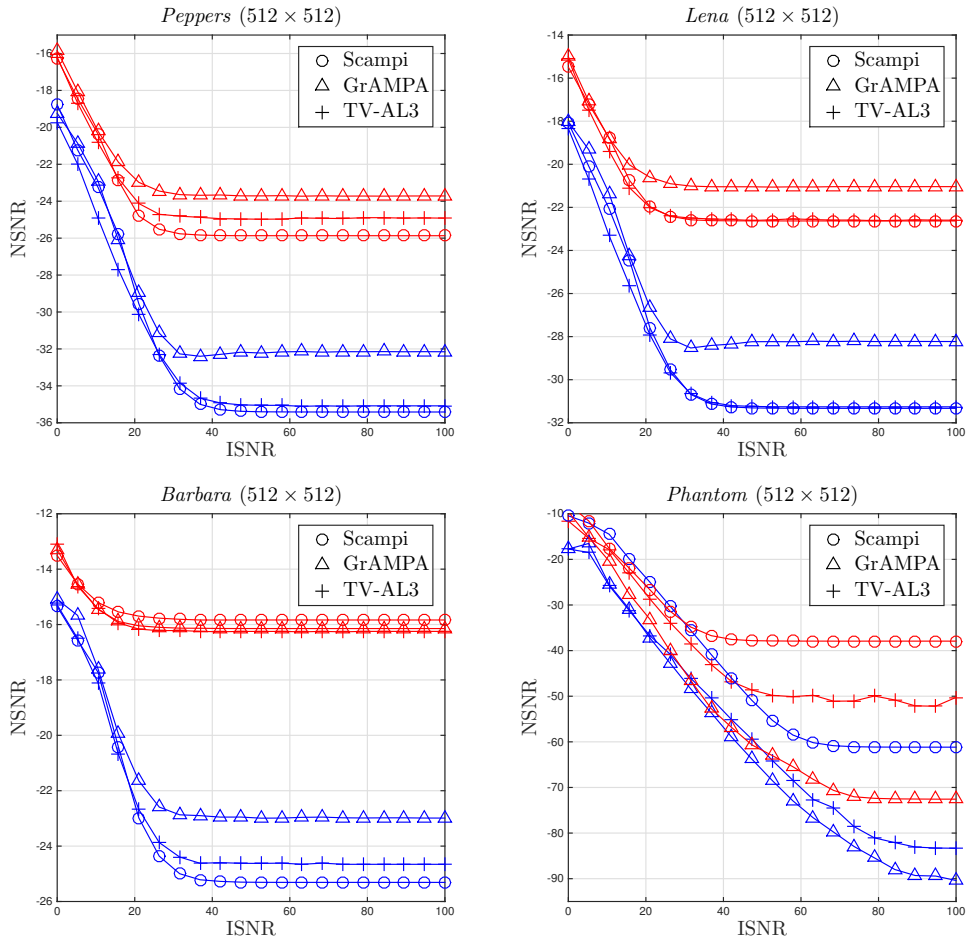

Figure 4. Comparison of optimal reconstruction accuracy as a function of the ISNR for Scampi, GrAMPA, and TV-AL3. All test points correspond to a search over 30 test-points for tuning the regularization parameter of TV-AL3 or the SNIPE ω parameter for GrAMPA and Scampi. Two measurement rates are tested for each method on each 512×512 image, $\alpha = 0.5$ (blue) and $\alpha = 0.1$ (red).

all tested natural images, TV-AL3 performs better than GrAMPA in terms of NSNR. Only for *Phantom*, whose gradient is very well modeled by the SNIPE prior, does GrAMPA outperform TV-AL3. However, one should note that the tuning of the free parameters of TV-AL3 is more problematic *a posteriori* than GrAMPA, as the domain of the optimal regularization term can be quite large, depending on the problem instance, while the optimal value of ω is generally around 0 or 1. But, these maximally best-case performance results imply that the convex treatment of TV is superior to the probabilistic model used by GrAMPA in the case of natural images.

However, the model used by Scampi, along with the dynamic estimation of both the measurement noise variance and constraint slack on the dual variables, provides best-case performance equal to or beyond TV-AL3's ones for most test cases, notably on *Peppers* and *Lena*, and at high measurement rates for an edge-heavy image such as *Barbara*. Finally, as we pointed out earlier, we see the expected subpar performance of Scampi for the cartoon *Phantom* image. However, we did not test the fixed $T = 0$ case for this image, which may yet improve the performance of Scampi for such images.

We also show on Fig. 5 some reconstruction results obtained with Scampi and GrAMPA for

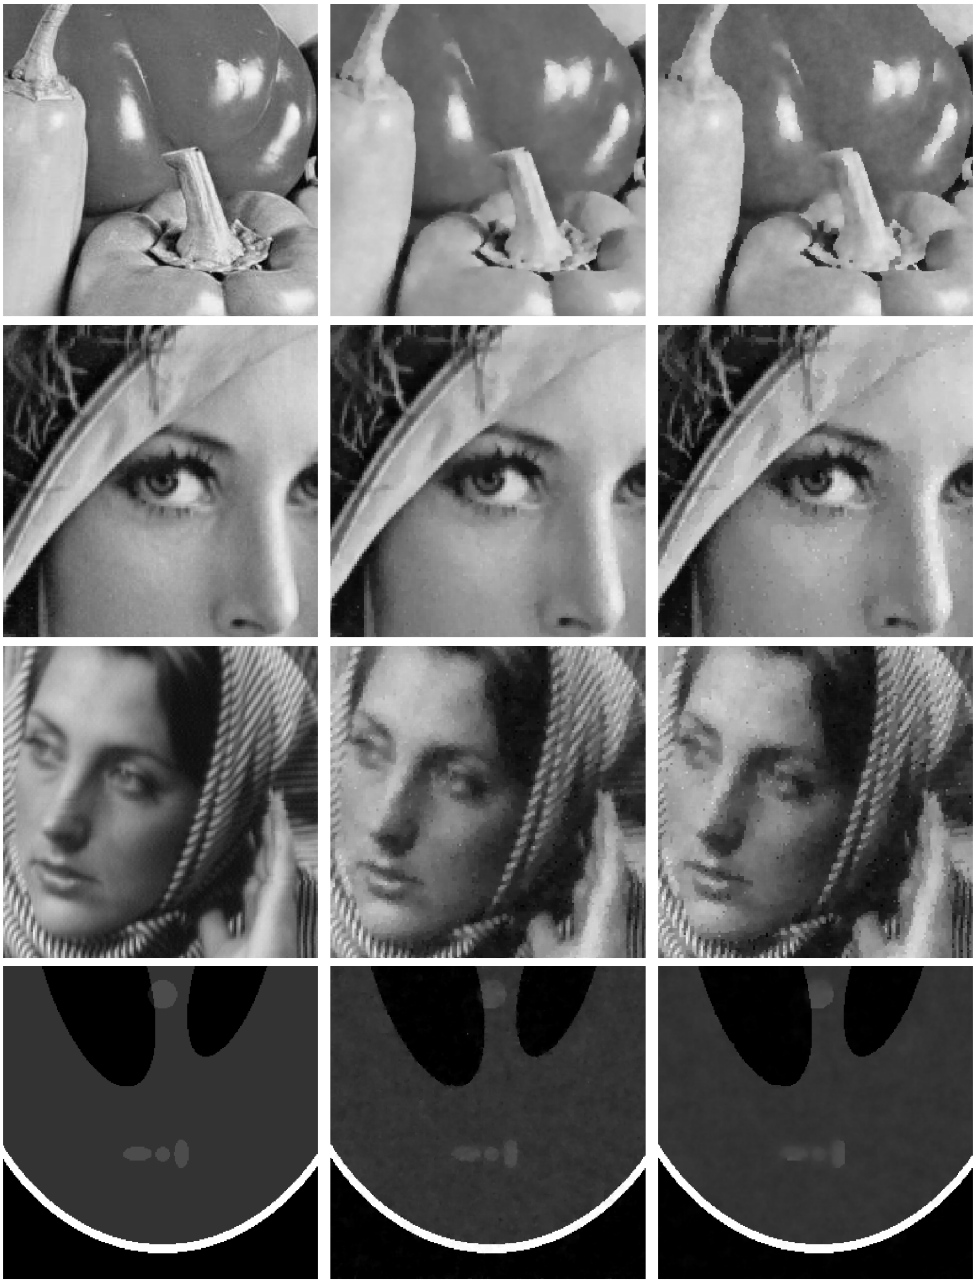

Figure 5. Visual comparison for some 512×512 images: a zoom on the original image (left) is compared with the Scampi (center) and GrAMPA (right) reconstructions for different settings, all with a fixed SNIPE prior parameter $\omega = 0$. From top to bottom are *peppers* at $\alpha = 0.1$ for an ISNR of 20 dB, *Lena* at $\alpha = 0.5$ for an ISNR of 40 dB, *Barbara* at $\alpha = 0.5$ for an ISNR of 30 dB and *phantom* at $\alpha = 0.1$ for an ISNR of 20 dB.

different settings, without optimization of ω which is set to 0 in all cases. As expected, the images obtained with Scampi are smoother and match better the three first natural images. In the last case of *phantom*, the gain in NSNR with GrAMPA with respect to Scampi is due to the large constant areas, but we notice that even in this case, Scampi gets a better resolution at the edges.

4. Conclusion

We have presented a new probabilistic AMP-based reconstruction algorithm for compressive imaging based on the cosparse analysis model with dual variables. This approach is similar to a finite temperature version of the GrAMPA algorithm, which makes it particularly adapted to natural images. Indeed, for this class of signals, the proposed algorithm *Scampi* reaches state -of-the-art perfomance, overcoming the previous best probabilistic methods. Furthermore, it is highly robust to noise channel uncertainty and does not require near-perfect knowledge to reach optimal performance. However, our experiments show that for cartoon images, consisting of only constant regions separated by instantaneous transitions, the zero temperature model does a better job of enforcing the piecewise-continuous prior on the image.

We have also shown through intensive numerical experiments that when we push the convex optimization TV-AL3 algorithm to its limits via oracle tuning parameter selection, it gives results very close to the proposed, easier to tune, approach. The fundamental reason for this empirical finding is yet to be understood.

Acknowledgments

This work has been supported in part by the ERC under the European Union's 7th Framework Programme Grant Agreement 307087-SPARCS, by the French Ministère de la défense/DGA and by the Swiss National Science Foundation grant number 200021-156672.

References

[1] Candès E, Romberg J and Tao T 2006 *Communications on Pure and Applied Mathematics* **59** 1207–1223
[2] Candès E J and Tao T 2006 *IEEE Trans. Inform. Theory* **52** 5406
[3] Candès E, Romberg J and Tao T 2006 *IEEE Trans. Inform. Theory* **52** 489–509
[4] Donoho D L 2006 *IEEE Trans. Inform. Theory* **52** 1289
[5] Duarte M F, Davenport M A, Takhar D, Laska J N, Sun T, Kelly K F and Baraniuk R G 2008 *Signal Processing Magazine* **25** 83–91
[6] Nam S, Davies M, Elad M and Gribonval R 2013 *Applied and Computational Harmonic Analysis* **34** 30 – 56
[7] Tibshirani R 1996 *Journal of the Royal Statistical Society, Series B* **58** 267–288
[8] Chambolle A and Pock T 2011 *Journal of Mathematical Imaging and Vision* **40** 120–145
[9] Li C 2009 *An Efficient Algorithm for Total Variation Regularization with Applications to the Single Pixel Camera and Compressive Sensing* Master's thesis Rice University
[10] Donoho D, Johnstone I and Montanari A 2013 *IEEE Trans. Inform. Theory* **59**
[11] Metzler C A, Maleki A and Baraniuk R G 2014 *arXiv Preprint* **cs.IT/1406.4175**
[12] Tan J, Ma Y and Baron D 2015 *IEEE Trans. Sig. Proc.* **63** 2085–2092
[13] Kang J, Jung H, Lee H N and Kim K 2014 *arXiv Preprint* **cs.IT/1408.3930**
[14] Rangan S 2011 Generalized approximate message passing for estimation with random linear mixing *Proc. IEEE Int. Symp. on Inform. Theory* pp 2168 –2172
[15] Borgerding M A, Schniter P and Rangan S 2013 *arXiv Preprint* **cs.IT/1312.3968**
[16] Donoho D L, Maleki A and Montanari A 2009 *Proc. Natl. Acad. Sci.* **106** 18914–18919
[17] Bayati M and Montanari A 2011 *IEEE Transactions on Information Theory* **57** 764 –785
[18] Donoho D L, Maleki A and Montanari A 2010 Message passing algorithms for compressed sensing: I. motivation and construction *Proc. IEEE Inform. Theory Workshop* pp 1–5
[19] Krzakala F, Mézard M, Sausset F, Sun Y F and Zdeborová L 2012 *Physical Review X* **2** 021005
[20] Krzakala F, Mézard M, Sausset F, Sun Y and Zdeborová L 2012 *Journal of Statistical Mechanics: Theory and Experiment* **2012** P08009
[21] Javanmard A and Montanari A 2012 Subsampling at information theoretically optimal rates *Proc. IEEE Int. Symp. on Inform. Theory* pp 2431–2435
[22] Barbier J 2015 *ArXiv e-prints* **cs.IT/1511.01650**
[23] Krzakala F, Manoel A, Tramel E W and Zdeborová L 2014 Variational free energies for compressed sensing *Proc. IEEE Int. Symp. on Inform. Theory* pp 1499–1503
[24] Barbier J, Schulke C and Krzakala F 2015 *Journal of Statistical Mechanics: Theory and Experiment* **2015** P05013